


HDF5



Quincey Koziol

koziol@lbl.gov

NERSC Data and Analytics Services

March 21, 2016

NERSC User Group Meeting

What is HDF5?


- HDF5 == Hierarchical Data Format, v5
- A **data model**
 - Structures for data organization and specification
- Open source **software**
 - Translates data model objects to file
- Open **file format**
 - Designed for high volume or complex data



HDF5 is designed ...

- for high volume and/or complex data
- for every size and type of system (portable)
- for flexible, efficient storage and I/O
- to enable applications to evolve in their use of HDF5 and to accommodate new models
- to support long-term data preservation

HDF5 Data Model

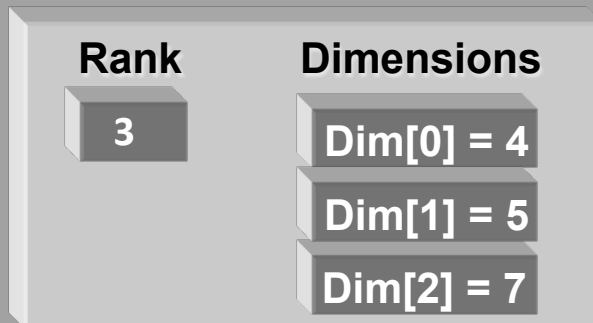


HDF5 Dataset


HDF5 Datatype

Integer: 32-bit, LE

HDF5 Dataspace



Specifications for single data element and array dimensions



Multi-dimensional array of identically typed data elements

- HDF5 datasets **organize and contain** data elements.
 - HDF5 datatype describes individual data elements.
 - HDF5 dataspace describes the logical layout of the data elements.

HDF5 Dataspace

Two roles:

Dataspace contains spatial information

- Rank and dimensions
- Permanent part of dataset definition



Rank = 2

Dimensions = 4x6

Partial I/O: Dataspace describes application's data buffer and data elements participating in I/O





Rank = 1

Dimension = 10

HDF5 Datatypes

- Describe individual data elements in an HDF5 dataset
- Wide range of datatypes supported
 - Integer
 - Float
 - Enum
 - Array
 - User-defined (e.g., 13-bit integer)
 - Variable-length types (e.g., strings, vectors)
 - Compound (similar to C structs)
 - More ...

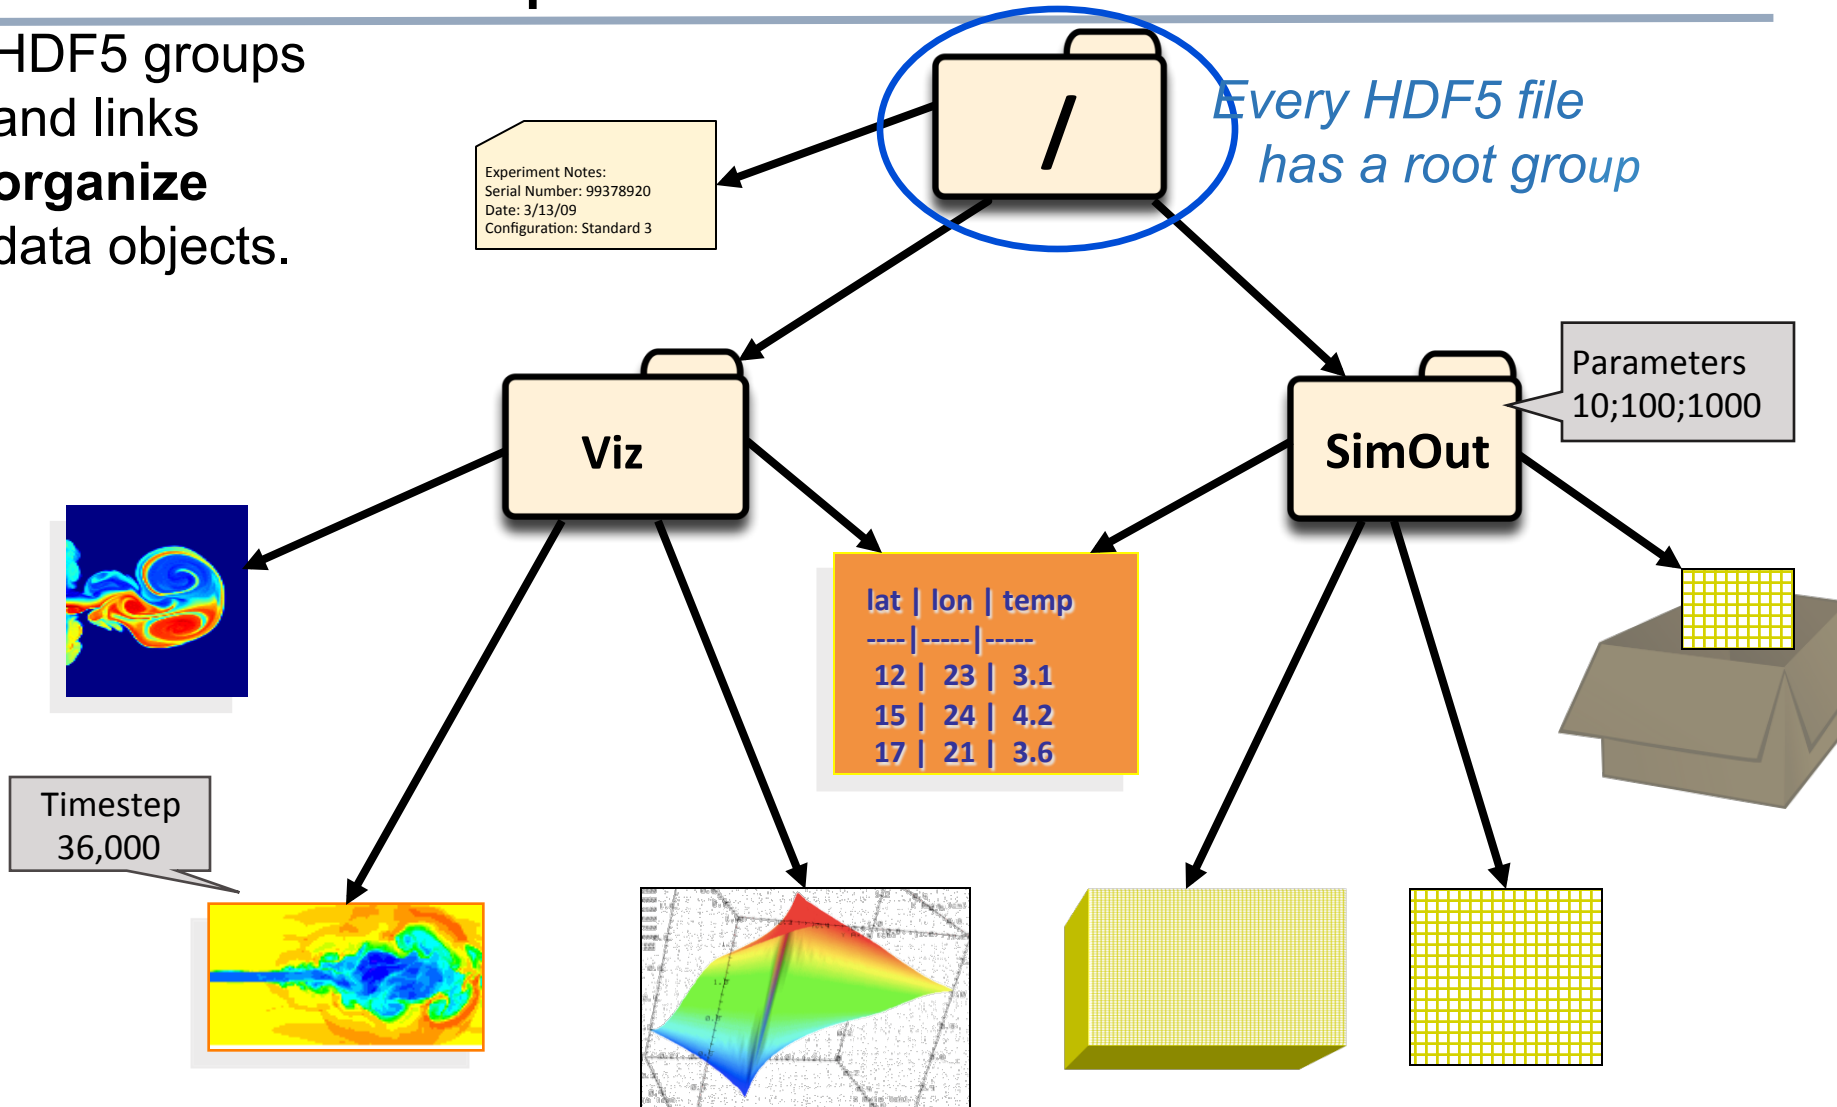
HDF5 Dataset with Compound Datatype



Dataspace: Rank = 2
Dimensions = 5 x 3

HDF5 Groups and Links

HDF5 groups and links **organize** data objects.



HDF5 Attributes

- Typically contain user metadata
- Have a name and a value
- Attributes “decorate” HDF5 objects
- Value is described by a datatype and a dataspace
- Analogous to a dataset, but do not support partial I/O operations; nor can they be compressed or extended

HDF5 Software

HDF5 home page: <http://hdfgroup.org/HDF5/>

- Latest release: HDF5 1.8.16 (1.10.0 coming April, 2016)

HDF5 source code:

- Written in C, and includes optional C++, Fortran 90 APIs, and High Level APIs
- Contains command-line utilities (h5dump, h5repack, h5diff, ..) and compile scripts

HDF5 pre-built binaries:

- When possible, include C, C++, F90, and High Level libraries. Check ./lib/libhdf5.settings file.
- Built with and require the SZIP and ZLIB external libraries

The HDF5 API

- C, Fortran, Java, C++, and .NET bindings
- IDL, MATLAB, R, Python (h5py, PyTables)
- C routines begin with prefix `H5?`
 ? is a character corresponding to the type of object
 the function acts on

Example Functions:

H5D : Dataset interface *e.g.*, **H5Dread**

H5F : File interface *e.g.*, **H5Fopen**

H5S : dataSpace interface *e.g.*, **H5Sclose**

The HDF5 API

- For flexibility, the API is extensive
 - 300+ functions
- This can be daunting... but there is hope
 - A few functions can do a lot
 - Start simple
 - Build up knowledge as more features are needed



Victorinox
Swiss Army
Cybertool 34



General Programming Paradigm

- Object is opened or created
 - Object is accessed, possibly many times
 - Object is closed
-
- Properties of object are optionally defined
 - Creation properties (e.g., use chunking storage)
 - Access properties

Basic Functions

H5**F**create (H5**F**open)

create (open) File

H5**S**create_simple/H5**S**create

create dataSpace

H5**D**create (H5**D**open)

create (open) Dataset

H5**D**read, H5**D**write

access Dataset

H5**D**close

close Dataset

H5**S**close

close dataSpace

H5**F**close

close File

Other Common Functions

Data S paces:	H5Sselect_hyperslab (Partial I/O) H5Sselect_elements (Partial I/O) H5Dget_space
Data T ypes:	H5Tcreate, H5Tcommit, H5Tclose H5Tequal, H5Tget_native_type
G roups:	H5Gcreate, H5Gopen, H5Gclose
A tttributes:	H5Acreate, H5Aopen_name, H5Aclose, H5Aread, H5Awrite
P roperty lists:	H5Pcreate, H5Pclose H5Pset_chunk, H5Pset_deflate

Enabling Collective Parallel I/O with HDF5

```
/* Set up file access property list w/parallel I/O access */
fa_plist_id = H5Pcreate(H5P_FILE_ACCESS);
H5Pset_fapl_mpio(fa_plist_id, comm, info);

/* Create a new file collectively */
file_id = H5Fcreate(filename, H5F_ACC_TRUNC,
    H5P_DEFAULT, fa_plist_id);


/* <omitted data decomposition for brevity> */

/* Set up data transfer property list w/collective MPI-IO */
dx_plist_id = H5Pcreate(H5P_DATASET_XFER);
H5Pset_dxpl_mpio(dx_plist_id, H5FD_MPIO_COLLECTIVE);

/* Write data elements to the dataset */
status = H5Dwrite(dset_id, H5T_NATIVE_INT,
    memspace, filespace, dx_plist_id, data);
```

H5Dwrite vs. H5Dwrite_multi

Contiguous floating-point datasets



Useful Tools For New Users

h5dump:

Tool to “dump” or display contents of HDF5 files

h5cc, h5c++, h5fc:

Scripts to compile applications

HDFView:

Java browser to view HDF5 files

<http://www.hdfgroup.org/products/java/hdfview/>

HDF5 Examples (C, Fortran, Java, Python, Matlab)

<http://www.hdfgroup.org/ftp/HDF5/examples/>

Questions, Comments, Feedback

?